Taming Weak Memory Models

Sizhuo Zhang
szzhang@mit.edu

Arvind
arvind@csail.mit.edu

Muralidaran Vijayaraghavan
vmurali@csail.mit.edu

Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology

arXiv:1606.05416v1 [cs.PL] 17 Jun 2016

ABSTRACT

Speculative techniques in microarchitectures relax various dependencies in programs, which contributes to the complexity of (weak) memory models. We show using WMM, a new weak memory model, that the model becomes simpler if it includes load-value speculation and thus, does not enforce any dependency! However, in the absence of good value-prediction techniques, a programmer may end up paying a price for the extra fences. Thus, we also present WMM-D, which enforces the dependencies captured by the current microarchitectures. WMM-D is still much simpler than other existing models. We also show that non-atomic multi-copy stores[1] arise as a result of sharing write-through caches. We think restricting microarchitectures to write-back caches (and thus simpler weak memory models) will not incur any performance penalty. Nevertheless, we present WMM-S, another extension to WMM, which could model the effects of non-atomic multi-copy stores.

WMM, WMM-D, and WMM-S are all defined using *Instantaneous Instruction Execution* (I^2E), a new way of describing memory models without explicit reordering or speculative execution.

1. INTRODUCTION

Architects have often made changes for performance in microarchitectures, which end up affecting the memory model of multiprocessors in subtle ways. The problem is serious because the official definitions of memory models, such as Power and ARM, are underspecified in the company documents [2, 3]. In order to resolve the ambiguity in the official documents, researchers have used additional empirical evidence and developed abstract machines and axioms to capture the behaviors of these models precisely. This situation is not satisfactory either, because these formal definitions use complex abstract machines and specify a large number of axioms to capture all the corner cases [4, 5, 6]. Additionally, most architects are intimidated by these formalizations.

If the behaviors to be modeled are complicated, then there is no reason to believe that the model definition can be simple. Indeed we believe this is the case with Power and ARM models, because they must capture the non-atomic multi-copy stores. Since these types of stores arise as a consequence of sharing write-through caches, one way to simplify the behaviors is to use write-back caches exclusively. Another source of complexity comes from the fact that speculative execution selectively relaxes program dependencies. It is complicated to model precisely the forbidden behaviors that are caused by specific dependencies.

Sometimes memory models have been defined *intuitively but incorrectly*. We will show that the RC definition [7] is imprecise both in dependency ordering and the values returned by loads, and that it is not easy to fix the model without making it too restrictive. We will also show that the RMO definition [8] is fairly precise, but rules out certain architectural optimizations which are commonly believed to be permitted.

We want a memory model that simultaneously satisfies the following three criteria:

1. Model definition is simple and precise;
2. Model allows efficient hardware implementations;
3. High-level language (HLL) primitives can be efficiently mapped to memory instructions and fences in the model.

By *simple* definition, we mean an *Instantaneous Instruction Execution* (I^2E) description of the instruction set. All abstract machines in the I^2E framework consist of a monolithic multi-ported memory and n atomic processors. The atomic processor executes each instruction instantaneously so the architecture state is by definition always up-to-date. The instruction reordering and dependencies in specific models can be captured by including different types of buffers between each processor and its port in the memory. I^2E descriptions are considerably simpler than other model definitions which either use reordering axioms or operational models that execute instructions partially.

The third criterion is important because people prefer to implement concurrent programs using primitives in HLLs instead of architecture-specific fences in assembly, and this requires a systematic mapping from HLL semantics to weak memory models.

None of the current architectural memory models satisfy all the three criteria. SC and TSO have simple and precise descriptions, but leave a lot on the table as far as performance is concerned. As already mentioned, weak memory models like Power, ARM, RC, and RMO do not satisfy the first criterion. The situation is fur-

[1] Also known as non-multi-copy-atomic stores [1].

ther complicated by the fact that there are attempts to design HLLs and extend weak models so that the two would be more compatible.

The major contributions of this paper are:

1. Identification of architectural features that cause non-atomic multi-copy stores in Power and ARM models;
2. The first characterization of the problems with RC and RMO definitions;
3. I^2E, a new framework for describing memory models;
4. WMM, an I^2E weak memory model which satisfies the three criteria of goodness given earlier, and its implementation with speculations on all dependencies;
5. WMM-D, an extension of WMM to capture data-dependency ordering precisely while still using I^2E.
6. WMM-S, an extension of WMM to model non-atomic multi-copy stores.

Paper organization: Section 2 presents the related work. We analyze the source of complexity in Power and ARM memory models in Section 3. We identify the problems in the definitions of RC and RMO in Section 4. In Section 5, we propose a new framework called I^2E to define memory models. We use I^2E to give simple abstract machines of SC, TSO and PSO. We define WMM using I^2E in Section 6, which also includes its implementation, litmus tests and a compilation scheme from C++. We define WMM-D to capture data dependency in Section 7. In Section 8, we extend WMM to WMM-S to model non-atomic multi-copy stores. Our conclusions are presented in Section 9.

2. RELATED WORK

SC [9] is the most intuitive memory model, but naive implementations of SC suffer from poor performance. Gharachorloo et al. proposed load speculation and store prefetch to enhance the performance of SC [10]. Over the years, researchers have proposed more aggressive techniques to preserve SC [11, 12, 13, 14, 15, 16, 17, 18, 19]. Perhaps because of their hardware complexity, the adoption of these techniques in commercial microprocessor has been limited. Instead the manufactures and researchers have chosen to present weaker memory model interfaces, e.g. TSO [20], PSO [8], RMO [8], Processor Consistency [21], Weak Consistency (WC) [22], RC [7], CRF [23], Power [2] and ARM [3]. The tutorials by Adve et al. [24] and by Maranget et al. [25] provide relationships among some of these models.

The lack of clarity in the definitions of Power and ARM memory models in their respective company documents has led some researchers to empirically determine allowed/disallowed behaviors [4, 5, 26, 6]. Based on such observations, in the last several years, both *axiomatic* models and *operational* models have been developed which are compatible with each other [27, 28, 29, 5, 26, 4, 30, 31, 6]. However, these models are quite complicated; for example, the Power axiomatic model has 10 relations, 4 types of events per instruction, and 13 complex axioms [5], some of which have been added over time to explain specific behaviors [27, 28, 32, 5]. The abstract machines used to describe Power and ARM operationally are also quite complicated, because they require the user to think in terms of partially executed instructions [4, 30]. In particular, the processor sub-model incorporates ROB operations, speculations, instruction replay on speculation failures, etc., explicitly, which are needed to explain the enforcement of specific dependency (i.e. data dependency). We present an I^2E model WMM-D in Section 7 that captures data dependency and sidesteps all these complications. Another source of complexity is in the memory sub-model, which we explain in Section 3.

Adve et al. defined Data-Race-Free-0 (DRF0), a class of programs where shared variables are protected by locks, and proposed that DRF0 programs should behave as SC [33]. However, architectural memory models must also define the behaviors of non-DRF0 programs.

A large amount of research has also been devoted to specifying the memory models of HLLs: C++ [34, 35, 36], Java [37, 38, 39], etc. We will provide compilation schemes from C++, a widely-used HLL, to the WMM and WMM-D models presented in this paper.

Arvind and Maessen have specified precise conditions for preserving store atomicity in program execution even when instructions can be reordered [40]. In contrast, the WMM and WMM-D models presented in this paper do not insist on store atomicity at the program level.

Recently, Lustig et al. have used Memory Ordering Specification Tables (MOSTs) to describe memory models, and proposed a hardware scheme which dynamically converts programs across memory models described in MOSTs [1]. MOST specifies the ordering strength (e.g. locally ordered, multi-copy atomic) of two instructions from the same thread under different conditions (e.g. data dependency, control dependency). It is not clear to us what events in the program execution are being (re)ordered by MOST. It is also unclear regarding which value a load returns given a legal order of events.

3. NON-ATOMIC MULTI-COPY STORES: AN AVOIDABLE COMPLICATION

Stores in TSO (Intel) are known as *multi-copy atomic*, because a store first becomes visible to the local processor and then later to all other processors simultaneously. In contrast, Stores in Power and ARM processors are *non-atomic multi-copy*, that is, a store may become visible to different processors at different times. This is caused by the memory system, which allows a store S from processor Pi to be observed by Pj before S has finished all coherence transactions. There are two root causes for this behavior. If multiple threads on a single core, i.e. as in Simultaneous Multithreading (SMT), share a store buffer, then a store by any of these threads may be seen by all these threads before other processors. This non-atomicity can be avoided by keeping the stores of threads separate by tagging them with thread IDs in the store buffer. If such tagged shared store buffers are combined with write-back caches, all threads except the one which issued the store cannot observe this store until the store is committed to L1. On the other hand, if

multiple threads share a write-through cache (typically L1), then these threads can see a store by any of these threads before the store reaches coherence, i.e. become globally visible, making it non-atomic. Unlike the case of shared store buffer, it is infeasible to distinguish between stores by different threads in the write-through cache. Even without SMT, the non-atomicity problem will persist if the shared L2 is write-through.

Sorin et al. have identified that it is hard to implement TSO with shared write-through caches [41, page 180]. The above analysis also matches their understanding. Later we will show that a weak memory model with multi-copy atomic stores, is considerably simpler than Power and ARM models. In case one really wants to model a non-atomic multi-copy memory system, we also present such a model using I^{2}E in Section 8.

4. INCORRECT MEMORY MODEL DEFINITIONS

It is generally believed that RC and RMO are well-defined memory models. We first show that the RC definition [7], in fact, is underspecified in the sense that it does not precisely define the values returned by loads, and the ordering of events in case of dependent instructions. We will further show that after several attempts to resolve the ambiguity, the resulting model is no longer weak enough. RMO definition [8], on the other hand, has a precise axiomatic description, but it fails to match an architect's intuitive understanding of the model. (Readers could skip this section without losing continuity, though later sections will refer to program examples used here).

4.1 RC

4.1.1 Original RC Definition [7]

A memory access in RC has an "*issue*" event followed by n "*being performed with respect to (w.r.t.) processor* i" ($i = 1 \ldots n$) events. The order of "issue" and "being performed w.r.t." events constrain the load values in two ways (see Definition 2.1 in [7]):

- *RC-LdVal-1*: If load L is performed w.r.t. Pi before store S is issued by Pi, L cannot read from S.
- *RC-LdVal-2*: If store S to address a is performed w.r.t. Pi before load L to a is issued by Pi, L must read from S or another store S' which is performed w.r.t. Pi after S.

When a memory access is performed w.r.t. all processors, the access is *performed*. Notice that Definition 2.1 of [7] is ambiguous about the load result in case that L is issued before S is performed while L is performed after S is performed. Thus, our attempted interpretation, RC-LdVal-2, only constrains the load result in case L is issued after S is performed.

RC classifies memory accesses into two categories: *ordinary* ones and *special* ones. Special accesses are further partitioned into three types: *acquire* (acq), *release* (rel) and *non-synchronization*. RC requires ordinary accesses to be *performed* before or after release and acquire accesses, respectively (see Condition 3.1 in [7]). In addition, RC places the following three constraints (see the last paragraph of Section 2 of [7]):

- *RC-Dependency*: "Uniprocessor control and data dependences are respected".
- *RC-Coherence*: "All writes to the same location are serialized in some order and are performed in that order w.r.t. any processor". The last store in that order gives the final memory value for that location.
- *RC-No-Deadlock*: "Accesses that occur previously in program order eventually get performed".

A legal execution is a total order of all events which satisfies all the constraints.

4.1.2 Ambiguity in RC Definition

We explain the ambiguity using examples. All examples in this paper assume that all memory locations are initialized to 0.

Load results: Consider the program in Figure 1, which is a common usage of acquire-release pair to communicate data. The behavior in the figure is allowed by the above definition, which is unexpected. Consider the following event order: I_1 to I_4 are issued one by one, and then I_1 to I_4 are performed one by one. Since there is no dependency in the program, events can be issued at the beginning, and the above order is legal. Since both I_1 and I_4 are issued before any of them is performed w.r.t. any processor, neither RC-LdVal-1 nor RC-LdVal-2 could constrain I_4 to only read from I_1. Thus I_4 is allowed to read the initial value 0.

Proc. P1	Proc. P2
I_1 : St a 1	$I_3 : r_1 = \mathsf{Ld}_{acq}\ b$
I_2 : St_{rel} b 1	$I_4 : r_2 =$ Ld a
RC allows: $r_1 = 1,\ r_2 = 0$	

Figure 1: Message passing in RC

Proc. P1	Proc. P2
$I_1 : r_1 =$ Ld a	$I_3 : r_2 =$ Ld b
I_2 : St b r_1	I_4 : St a r_2
RC allows: $r_1 = r_2 = 42$	

Figure 2: Thin-air read in RC

Ambiguous dependency constraint: There are two interpretations of the RC-Dependency constraint. Suppose it means that if an access B depends on another access A, then B cannot be issued before A is performed. Obviously with this interpretation, no speculative execution is possible.

Another interpretation is that dependency does not enforce any ordering of events, as long as the final result satisfies the program logic. This interpretation is so relaxed that it allows the "thin-air read" behavior as shown in Figure 2. This behavior may incur security problems, and is explicitly forbidden by C++ [34] for relaxed atomic loads and stores. Thus the compilation from C++ to RC will be inefficient.

RC definition is also unclear about how two loads to the same address on the same processor affect the ordering of events.

4.1.3 Attempts to Fix RC

Since RC is an extension for WC, suppose we constrain the load values by borrowing Dubois' et al. [22] WC solution:

- *RC-Fix-WC-Ld-Value*: The result of a load on Pi should be the value given by the latest store (for the same address) performed w.r.t. Pi.

With RC-Fix-WC-Ld-Value, when a load gets its result by reading from a store, the load is automatically performed w.r.t. all processors. (It is easy to see that RC-Fix-WC-Ld-Value implies RC-LdVal-1 and RC-LdVal-2, so we do not need these two earlier conditions).

We can fix the dependency constraint to allow speculation while avoiding "thin-air" read as follows:

- *RC-Fix-Dependency*: A store S that depends on a load L should not be performed w.r.t. any other processor before L is performed.

Now consider the Write-Write Causality (WWC) program in Figure 3. According to the RC definition after fixes, I_1 will be performed before I_5 w.r.t. P2, thus making it impossible for the $m[a]$ to be 2 according to RC-Coherence. However, most architects believe that RC can be implemented with a non-atomic multi-copy memory system (e.g. P1 and P2 shares a write-through L1), which allows this behavior.

Furthermore, if we move I_1 from P1 to P2 and place I_1 right before I_2 in Figure 3, the behavior in the figure is still disallowed by RC for the same reason. However, the behavior is observable in implementations, in which P2 could locally forward the data of I_1 to I_2. And such implementations are also believed to be permitted by RC. Thus, the gap between the intuitive understanding of RC and its precise definition remains. It is probably possible to come up with a much more complex fix for RC, but then it will not be a simple model anymore.

Proc. P1	Proc. P2	Proc. P3
I_1 : St a 2	$I_2 : r_1 = \mathsf{Ld}\ a$ I_3 : St b $(r_1 - 1)$	$I_4 : r_2 = \mathsf{Ld}\ b$ I_5 : St a r_2
RC forbids: $r_1 = 2,\ r_2 = 1,\ m[a] = 2$		

Figure 3: WWC (non-atomic multi-copy stores) in RC

4.2 RMO

There is also a gap between the definition of RMO [8, Appendix D] and the optimizations which are expected to be allowed in its implementation. RMO forbids the behavior shown in Figure 4 because I_9 transitively depends on I_4 according to the definition, and I_9 is ordered after I_4 in memory order. However, this behavior is possible in hardware with speculative load execution and store forwarding, i.e. I_7 first speculatively bypasses from I_6, and then I_9 speculatively executes to get 0. Most architects will not be willing to give up on these two optimizations.

Besides the above problem, RMO permits the reordering of two loads to the same address, i.e. it allows the non-SC behavior in the Coherent Read-Read (CoRR) example in Figure 5. Although this is not wrong, it makes the compilation from C++ [34] inefficient, because such behavior is forbidden by C++ even for relaxed atomic loads and stores. One could fix this problem by adding the following axiom:

Proc. P1	Proc P2
I_1 : St a 1 I_2 : MEMBAR I_3 : St b 1	$I_4 : r_1 = \mathsf{Ld}\ b$ I_5 : **if**$(r_1 \neq 1)$ exit I_6 : St c 1 $I_7 : r_2 = \mathsf{Ld}\ c$ $I_8 : r_3 = a + r_2 - 1$ $I_9 : r_4 = \mathsf{Ld}\ r_3$
RMO forbids: $r_1 = 1,\ r_2 = 1$ $r_3 = a,\ r_4 = 0$	

Figure 4: Speculation in RMO

Proc. P1	Proc. P2
$I_1 : r_1 = \mathsf{Ld}\ a$ $I_2 : r_2 = \mathsf{Ld}\ a$	I_3 : St a 1
RMO allows: $r_1 = 1,$ $r_2 = 0$	

Figure 5: CoRR in RMO

- *RMO-Fix-Ld*: the program order of two loads to the same address must be preserved in memory order.

However this fix rules out an optimization implemented in Power and ARM processors (see Section 7.3).

5. DEFINING MEMORY MODELS USING I²E

We will model multiprocessor systems as shown in Figure 6 to define memory models. The state of the system with n processors is defined as $\langle ps, m\rangle$, where m is an n-ported *monolithic* memory which is connected to the n processors. $ps[i]$ $(i = 1 \ldots n)$ represents the state of the i^{th} processor. Each processor contains a register state s, which represents all architectural registers, including both the general purpose registers and special purpose registers, such as PC. Cloud represents additional state elements, e.g. a *store buffer*, that a specific memory model may use in its definition.

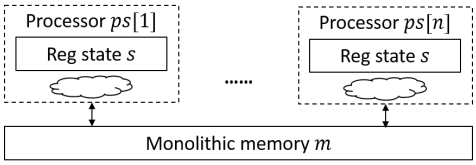

Figure 6: General model structure

5.1 Abstracting the Instruction Set

Memory model is always part of the ISA. However, we want our definitions of the memory models to be as generic as possible. For this reason, we introduce the concept of *decoded instruction set* (DIS). A *decoded instruction* contains all the information of an instruction after it has been decoded and has read all source registers. To begin with our DIS has the following three instructions.

- $\langle \mathsf{Nm}, dst, v\rangle$: instructions that do not access memory, such as ALU or branch instructions. It writes the computation result v into destination register dst.
- $\langle \mathsf{Ld}, a, dst\rangle$: a load that reads memory address a and updates the destination register dst.
- $\langle \mathsf{St}, a, v\rangle$: a store that writes value v to memory address a.

Later we will extend the DIS with fence instructions as needed. Next we explain how we get to decoded instructions from the source or raw instructions.

Instantaneous Instruction Execution (I^{2}E): To define memory models we restrict ourselves to I^{2}E models where each instruction is executed instantaneously and the register state of each processor is by definition always up-to-date. Therefore we can define the following two methods on each processor to manipulate the register state s:

- decode(): fetches the next raw instruction and returns the corresponding decoded instruction based on the current register state s.
- execute($dIns, ldRes$): updates the register state s (e.g. by writing destination registers and changing PC) according to the current decoded instruction $dIns$. A Ld requires a second argument $ldRes$ which should be the loaded value. For other instructions, the second argument can be set to don't care ("-").

I^{2}E cannot describe the semantics of a memory model in which the meaning of an instruction may depend upon a future store. Therefore all I^{2}E models we discuss do no permit stores to overtake loads.

5.2 Notations for Operational Semantics

The operational semantics is a set of *rules* that describe how the state of the abstract machine evolves as execution progresses. Each rule takes the following form:

$$\frac{predicates\ (based\ on\ the\ current\ state)}{actions\ (to\ modify\ the\ current\ state)}$$

If all *predicates* of a rule are *satisfied* then it can *fire* and atomically update model states according to the specified *actions*.

A predicate is either a *when* statement or a *pattern matching* statement. For example, when(b.empty()) means that the rule requires buffer b to be empty in order to fire. The *pattern matching* statement has the following form:

$$pattern = expression$$

For example, if we want to match the instruction returned by the decode() method to be a Nm instruction, we can write $\langle \mathsf{Nm}, dst, v\rangle = ps[i].\mathsf{decode}()$. Free variables dst and v will be assigned to appropriate values if the matching is successful. The matching identifier always begins with a capital letter, e.g. Nm, Ld, St, etc.

We use "$\Leftarrow$" to assign a new value to a state, and use semicolon ";" to separate statements written on the same line. If multiple rules can fire, then our semantic model selects any one of those rules *non-deterministically* to execute. The final outcome may depend on the choice of rule selection. To better understand the notation, we will give I^{2}E descriptions of three well-known memory models: SC, TSO and PSO (see Figures 7a and 7b).

5.3 SC Model

As shown in Figure 7a, SC does not require any special buffer. Figure 8 shows the operational semantics of SC. The three rules correspond to the *instantaneous execution* of the three types of decoded instructions. In each rule, the decode() method first fetches and decodes

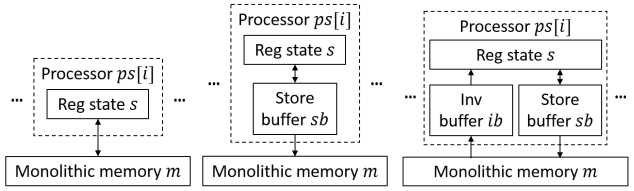

(a) SC (b) TSO/PSO (c) WMM

Figure 7: Structures of I^{2}E models

a new instruction, and then the instruction is immediately executed and committed. Loads and stores in SC directly access the monolithic memory.

SC-Nm rule (Nm execution).

$$\frac{\langle \mathsf{Nm}, dst, v\rangle = ps[i].\mathsf{decode}()}{ps[i].\mathsf{execute}(\langle \mathsf{Nm}, dst, v\rangle, \text{-})}$$

SC-Ld rule (Ld execution).

$$\frac{\langle \mathsf{Ld}, a, dst\rangle = ps[i].\mathsf{decode}()}{ps[i].\mathsf{execute}(\langle \mathsf{Ld}, a, dst\rangle, m[a])}$$

SC-St rule (St execution).

$$\frac{\langle \mathsf{St}, a, v\rangle = ps[i].\mathsf{decode}()}{ps[i].\mathsf{execute}(\langle \mathsf{St}, a, v\rangle, \text{-});\ m[a] \Leftarrow v}$$

Figure 8: SC operational semantics

5.4 TSO Model

Figure 7b shows the states and structure of TSO operational model proposed in [42, 43]. In addition to register state s, each processor now contains a store buffer sb. sb is an unbounded buffer of $\langle$address, value$\rangle$ pairs, each representing a pending store. The following methods are defined on sb:

- enq(a, v): enqueues the $\langle$address, value$\rangle$ pair $\langle a, v\rangle$ into sb.
- deq(): deletes the oldest store from sb, and returns its $\langle$address, value$\rangle$ pair.
- empty(): returns True when sb is empty.
- exist(a): returns True if address a is present in sb.
- youngest(a): returns the store data of the youngest store to address a in sb.

For methods that update the state and return a value, we use "$\leftarrow$" as in $\langle a, v\rangle \leftarrow sb.\mathsf{deq}()$, which assigns the return value of deq() to pair $\langle a, v\rangle$.

In order to enforce instruction ordering in accessing the newly added store buffer, we extend our instruction set with the memory fence instruction called Commit which flushes the local store buffer.

Figure 9 shows the operational semantics of TSO. The first four rules are instantaneous execution of four types of decoded instructions, while the fifth rule handles the interaction between the store buffer and monolithic memory. According to TSO-Ld, Ld a first tries to read the youngest store to address a in the local sb, and if sb does not contain a then it reads the monolithic memory. Buffering stores in sb essentially allows

a load to overtake a store. The Commit fence blocks until older stores are flushed from the store buffer. The buffer will eventually get empty as the consequence of repeated execution of TSO-DeqSb.

TSO-Nm rule (Nm execution).

$$\frac{\langle \mathsf{Nm}, dst, v\rangle = ps[i].\mathsf{decode}()}{ps[i].\mathsf{execute}(\langle \mathsf{Nm}, dst, v\rangle, \text{-})}$$

TSO-Ld rule (Ld execution).

$$\frac{\langle \mathsf{Ld}, a, dst\rangle = ps[i].\mathsf{decode}() \quad v = \textbf{if } ps[i].sb.\mathsf{exist}(a) \textbf{ then } ps[i].sb.\mathsf{youngest}(a) \textbf{ else } m[a]}{ps[i].\mathsf{execute}(\langle \mathsf{Ld}, a, dst\rangle, v)}$$

TSO-St rule (St execution).

$$\frac{\langle \mathsf{St}, a, v\rangle = ps[i].\mathsf{decode}()}{ps[i].\mathsf{execute}(\langle \mathsf{St}, a, v\rangle, \text{-});\ ps[i].sb.\mathsf{enq}(a, v)}$$

TSO-Com rule (Commit execution).

$$\frac{\langle \mathsf{Commit}\rangle = ps[i].\mathsf{decode}();\ \mathsf{when}(ps[i].sb.\mathsf{empty}())}{ps[i].\mathsf{execute}(\langle \mathsf{Commit}\rangle, \text{-})}$$

TSO-DeqSb rule (dequeue TSO store buffer).

$$\frac{\mathsf{when}(\neg ps[i].sb.\mathsf{empty}())}{\langle a, v\rangle \leftarrow ps[i].sb.\mathsf{deq}();\ m[a] \Leftarrow v}$$

PSO-DeqSb rule (dequeue PSO store buffer).

$$\frac{a = ps[i].sb.\mathsf{anyAddr}();\ \mathsf{when}(a \neq \epsilon)}{v \leftarrow ps[i].sb.\mathsf{rmOldest}(a);\ m[a] \Leftarrow v}$$

Figure 9: TSO/PSO operational semantics

5.4.1 PSO: Enabling Store-Store reordering

We extend TSO to PSO by allowing *sb* to commit the oldest store of any address to the monolithic memory, i.e. replacing the TSO-DeqSb rule by the PSO-DeqSb rule as shown in Figure 9. We use the following two methods, instead of deq(), to delete entries from *sb*.

- anyAddr(): returns any store address present in *sb*; or returns ϵ if *sb* is empty.
- rmOldest(a): deletes the oldest store to address a from *sb*, and returns its store data.

Thus we can dequeue stores for different addresses from the same store buffer out of order, i.e. reorder stores.

6. WMM MODEL

WMM allows Load-Load reordering in addition to the reorderings allowed by PSO. Since a reordered load may read a stale value, we introduce a conceptual device called *invalidation buffer*, *ib*, to each processor (see Figure 7c). *ib* is an unbounded buffer of $\langle$address, value$\rangle$ pairs, each representing a stale memory value for an address that can be observed by the processor. A stale value enters *ib* when some store buffer pushes a value to the monolithic memory. Similar to Commit fences for store buffers, we need a Reconcile memory fence to flush the local invalidation buffer.

The following methods are defined on *ib*:

- insert(a, v): inserts $\langle$address, value$\rangle$ pair $\langle a, v\rangle$ into *ib*.
- getRandom(a): returns a random value v for address a present in *ib*, and removes all values for a, which are inserted into *ib* before v, from *ib*.
- clear(): removes all contents from *ib* to make it empty.
- rmAddr(a): removes all (stale) values for address a from *ib*.

clear and rmAddr are called when ordering needs to be enforced, i.e. preventing younger loads from reading stale values.

6.1 Operational Semantics of WMM

Figure 10 shows the operational semantics of WMM. The first 7 rules are the instantaneous execution of decoded instructions, while the WMM-DeqSb rule removes the oldest store for any address (say a) from *sb* and commits it to the monolithic memory. WMM-DeqSb also inserts the original memory value into the *ib* of all other processors to allow Ld a in these processors to effectively get reordered with older instructions. However, this insertion in *ib* should not be done if the corresponding *sb* on that processor already has a store to a. This restriction is important, because if a processor has address a in its *sb*, then it can never see stale values for a. For the same reason, when a St a v is inserted into *sb*, we remove all values for a from the *ib* of the same processor.

WMM-Nm rule (Nm execution).

$$\frac{\langle \mathsf{Nm}, dst, v\rangle = ps[i].\mathsf{decode}()}{ps[i].\mathsf{execute}(\langle \mathsf{Nm}, dst, v\rangle, \text{-})}$$

WMM-LdSb rule (Ld execution: bypass from store).

$$\frac{\langle \mathsf{Ld}, a, dst\rangle = ps[i].\mathsf{decode}();\ \mathsf{when}(ps[i].sb.\mathsf{exist}(a))}{ps[i].\mathsf{execute}(\langle \mathsf{Ld}, a, dst\rangle, ps[i].sb.\mathsf{youngest}(a))}$$

WMM-LdMem rule (Ld execution: read memory).

$$\frac{\langle \mathsf{Ld}, a, dst\rangle = ps[i].\mathsf{decode}();\ \mathsf{when}(\neg ps[i].sb.\mathsf{exist}(a))}{ps[i].\mathsf{execute}(\langle \mathsf{Ld}, a, dst\rangle, m[a]);\ ps[i].ib.\mathsf{rmAddr}(a)}$$

WMM-LdIb rule (Ld execution: read stale value).

$$\frac{\langle \mathsf{Ld}, a, dst\rangle = ps[i].\mathsf{decode}() \quad \mathsf{when}(\neg ps[i].sb.\mathsf{exist}(a) \ \wedge\ ps[i].ib.\mathsf{exist}(a))}{v \leftarrow ps[i].ib.\mathsf{getRandom}(a);\ ps[i].\mathsf{execute}(\langle \mathsf{Ld}, a, dst\rangle, v)}$$

WMM-St rule (St execution).

$$\frac{\langle \mathsf{St}, a, v\rangle = ps[i].\mathsf{decode}()}{ps[i].\mathsf{execute}(\langle \mathsf{St}, a, v\rangle, \text{-});\ ps[i].sb.\mathsf{enq}(a, v);\ ps[i].ib.\mathsf{rmAddr}(a)}$$

WMM-Com rule (Commit execution).

$$\frac{\langle \mathsf{Commit}\rangle = ps[i].\mathsf{decode}();\ \mathsf{when}(ps[i].sb.\mathsf{empty}())}{ps[i].\mathsf{execute}(\langle \mathsf{Commit}\rangle, \text{-})}$$

WMM-Rec rule (Reconcile execution).

$$\frac{\langle \mathsf{Reconcile}\rangle = ps[i].\mathsf{decode}()}{ps[i].ib.\mathsf{clear}();\ ps[i].\mathsf{execute}(\langle \mathsf{Reconcile}\rangle, \text{-})}$$

WMM-DeqSb rule (dequeue store buffer).

$$\frac{a = ps[i].sb.\mathsf{anyAddr}();\ old = m[a];\ \mathsf{when}(a \neq \epsilon)}{\begin{array}{c} v \leftarrow ps[i].sb.\mathsf{rmOldest}(a);\ m[a] \Leftarrow v \\ \forall j \neq i.\ \textbf{if } \neg ps[j].sb.\mathsf{exist}(a) \textbf{ then } ps[j].ib.\mathsf{insert}(a, old)\end{array}}$$

Figure 10: WMM operational semantics

Load execution rules in Figure 10 correspond to three places from where a load can get its value. WMM-LdSb executes Ld a by reading from *sb*. If address a is not found in *sb*, then the load can read from the

monolithic memory (WMM-LdMem). However, in order to allow the load to read a stale value (to model load reordering), WMM-LdIb gets the value from ib. The model allows non-deterministic choice in the selection of WMM-LdMem and WMM-LdIb. To make this idea work, WMM-LdMem has to remove all values for a from ib, because these values are staler than the value in memory. Similarly, WMM-LdIb removes all the values for a, which are staler than the one read, from ib.

Synchronization instructions: Atomic read-modify-write (RMW) instructions can also be included in WMM. RMW should directly operate on the monolithic memory, so the rule to execute RMW is simply the combination of WMM-LdMem, WMM-St and WMM-DeqSb. One could also extend WMM to incorporate the load-linked/store-conditional pair in a similar way.

6.2 Litmus Tests for WMM

WMM executes instructions instantaneously and in order, but because of store buffers (sb) and invalidation buffers (ib), a processor can see the effect of loads and stores on some other processor in a different order than the program order on that processor. We explain the reorderings permitted and forbidden by the definition of WMM using well-known examples.

Fences for mutual exclusion: Figure 11 shows the kernel of Dekker's algorithm in WMM, which guarantees mutual exclusion by ensuring registers r_1 and r_2 cannot both be zero at the end. Fences I_2, I_3, I_6, I_7 are needed to keep this invariant. Suppose we remove a Reconcile fence I_3. Consider the scenario that all instructions on P2 execute first and I_8 gets 0. After that, all instruction on P1 execute and I_4 reads the stale value 0 from ib. It is as if I_4 overtakes I_1 and I_2. If we alternatively remove a Commit fence I_2, consider the case that all instructions on P1 execute first and I_4 gets 0. Let I_1 keep staying in the sb of P1. Meanwhile, we execute all instructions on P2, so I_8 will also get 0.

Proc. P1	Proc. P2
I_1 : St a 1	I_5 : St b 1
I_2 : Commit	I_6 : Commit
I_3 : Reconcile	I_7 : Reconcile
$I_4 : r_1 =$ Ld b	$I_8 : r_2 =$ Ld a
WMM forbids: $r_1 = 0, r_2 = 0$	

Figure 11: Dekker's algorithm in WMM

Proc. P1	Proc. P2
I_1 : St a 42	$I_4 : r_1 =$ Ld f
I_2 : Commit	I_5 : Reconcile
I_3 : St f 1	$I_6 : r_2 =$ Ld a
WMM forbids: $r_1 = 1, r_2 = 0$	

Figure 12: Message passing in WMM

Fences for message passing: Figure 12 shows a way of inter-processor communication in WMM. P1 writes data 42 to addresses a, and then signals P2 by setting a flag at address f to 1. P2 sees the new value of f and then reads the data. Fences I_2 and I_5 are needed to ensure that the data is correctly passed to P2. Without the Commit fence I_2, the data 42 may stay in the sb of P1 even after the flag has been set in the monolithic memory, and P2 may not see the new data. It is as if the two stores on P1 are reordered. Without the Reconcile fence I_5, P2 could see the stale value 0 from ib. It is as if the two loads on P2 are reordered.

No thin-air read: The Thin-air Read behavior in Figure 2 is impossible in WMM because of I^{2}E.

SC for a single address: WMM maintains SC for all accesses to a single memory location. This is because both the store buffer and invalidation buffer have the FIFO property for values of the same address. Therefore WMM does not appear to reorder loads to the same address, or stores to the same address. For example, the non-SC behavior in the CoRR example in Figure 5 is forbidden by WMM.

SC for well-synchronized programs: The critical sections in well-synchronized programs are all protected by locks. To maintain SC behaviors for such programs in WMM, we only need to add a Reconcile after acquiring the lock and a Commit before releasing the lock.

In summary, WMM can reorder stores to different addresses, and allows a load to overtake other loads (to different addresses), stores and Commit fences. WMM disallows a load to overtake any Reconcile fence.

6.3 Compiling C++ to WMM

C++ primitives [34] can be mapped to WMM instructions in an efficient way as shown in Table 1. For the purpose of comparison, we also include a mapping to Power which has been proven correct [44].

C++ operations	WMM instructions	Power instructions
Non-atomic Load	Ld	Ld
Load Relaxed	Ld	Ld
Load Consume	Ld; Reconcile	Ld
Load Acquire	Ld; Reconcile	Ld; cmp; bc; isync
Load SC	Commit; Reconcile; Ld; Reconcile	sync; Ld; cmp; bc; isync
Non-atomic Store	St	St
Store Relaxed	St	St
Store Release	Commit; St	lwsync; St
Store SC	Commit; St	sync; St

Table 1: Mapping C++ to WMM and Power

The Commit; Reconcile sequence in WMM is the same as a sync in Power, and Commit is similar to lwsync. The cmp; bc; isync sequence in Power can be viewed as a Load-Load fence, so it is similar to a Reconcile fence in WMM. In case of Store SC in C++, WMM uses a Commit while Power uses a sync, so WMM effectively saves one Reconcile. On the other hand, Power does not need any fence for Load Consume in C++, while WMM requires a Reconcile. Thus it is difficult to say whether one is more efficient than the other.

6.4 WMM Implementation

WMM can be implemented using modern OOO multiprocessors without any change in the microarchitecture, and even the most aggressive optimizations in ROB and cache cannot step beyond WMM. To demonstrate this, we first show how general OOO microarchitecture is abstracted by the WMM model, and then we discuss detailed hardware optimizations.

6.4.1 Correspondence of OOO to WMM

Let us consider an OOO microarchitecture (referred

below simply as OOO) with ROB, store buffer and a coherent write-back cache hierarchy. In OOO, instructions in ROB are committed in order, loads can be issued as soon as its address is known, and a store is enqueued into the store buffer only when the store commits (i.e. the entries in a store buffer cannot be killed). We show how WMM captures the behaviors of OOO by giving a correspondence between OOO and WMM.

Cache hierarchy: Although the write-back cache hierarchy may process many requests simultaneously and out of order, every request in the cache hierarchy completes by reading or writing an L1 cache line when it has sufficient permissions and the coherent memory value. The monolithic memory in WMM contains only such coherent values, and thus each access to L1 corresponds to an access to the monolithic memory in WMM. This monolithic memory abstraction of a coherent cache hierarchy has been proven by Vijayaraghavan et al. [45].

Store buffer: The state of the store buffer in OOO is represented by the *sb* in WMM. Entry into the store buffer when a store commits in OOO corresponds to the WMM-St rule. In OOO, the store buffer only issues the oldest store for some address to memory. The store is removed from the store buffer when the store updates L1. The removal from the store buffer exactly corresponds to the WMM-DeqSb rule.

ROB and eager loads: The commit of each instruction from ROB corresponds to the WMM rule that executes that instruction, and therefore the architectural state in both WMM and OOO must match at the time of commit. Early execution of a load L to address a with a return value v in OOO can be understood by considering where $\langle a, v\rangle$ resides in OOO when L *commits*. The WMM-LdSb and the WMM-LdMem rules cover the cases that $\langle a, v\rangle$ is, respectively, in the store buffer or the cache hierarchy when L commits. Otherwise $\langle a, v\rangle$ is no longer present in OOO at the time of load commit and must have been overwritten in memory. This case corresponds to using the WMM-DeqSb rule to inserts $\langle a, v\rangle$ into *ib*, and then using the WMM-LdIb rule to read v from *ib*.

Fences: Fences never go into the store buffer or memory. In OOO, Commit can commit from ROB only when the local store buffer is empty. Reconcile plays a different role; it stalls all younger loads unless the load can bypass from a store which is younger than the fence in ROB. The stall prevents younger loads from reading values that have become stale when the Reconcile commits. This corresponds to clearing *ib* in WMM.

In general, we can give a WMM execution for any OOO execution following the above correspondence. Each time the OOO execution commits an instruction I from ROB or removes a store S from store buffer, the coherent memory state, store buffers, and results of committed instructions in OOO are exactly the same as those in WMM when the WMM execution executes I or dequeues S from *sb*, respectively.

6.4.2 Aggressive Optimizations

Speculation: WMM does not enforce any dependency ordering, so the implementation can do all kinds of speculations, such as branch prediction, memory dependency prediction [46], and even load-value prediction [47, 48, 49, 50]. As a result, a load L can be issued as soon as we know its load address, which could even be computed from a predicted value. When all predictions related to L turn out to be correct, there is no need to check whether the value that L got earlier has become stale, because getting a stale value is captured by reading *ib* in WMM.

Consider the behavior in Figure 13. In an implementation with memory dependency prediction, P2 can predict that the store address of I_5 is not a, and execute I_6 early to get value 0. WMM allows this behavior because I_6 can read 0 from *ib*. Next consider the behavior in Figure 14. In an implementation with load-value prediction, P2 can predict the result of I_4 to be a and execute I_5 early to get value 0. When I_4 returns from memory later with value a, the prediction on I_4 turns out to be correct and the result of I_5 can be kept. WMM also allows this behavior because I_5 can read 0 from *ib*. Note that the behavior in Figure 14 is disallowed by RMO, Power and ARM if we change I_2 to MEMBAR (RMO fence), lwsync (Power fence) or dmb (ARM fence). Thus, load-value prediction cannot be directly used in RMO, Power and ARM processors.

Proc. P1	Proc. P2
I_1 : St a 1	$I_4 : r_1 =$ Ld b
I_2 : Commit	I_5 : St (r_1+c-1) 1
I_3 : St b 1	$I_6 : r_2 =$ Ld a
WMM allows: $r_1 = 1, r_2 = 0$	

Figure 13: Memory dependency prediction

Proc. P1	Proc. P2
I_1 : St a 1	$I_4 : r_1 =$ Ld b
I_2 : Commit	$I_5 : r_2 =$ Ld r_1
I_3 : St b a	
WMM allows: $r_1 = a, r_2 = 0$	

Figure 14: Load-value prediction

The only restriction on issuing loads comes from the fact that WMM does not appear to reorder loads for the same address. The implementation could execute such two loads out-of-order, but it must ensure that the return values of the two loads are from the same store. (Power and ARM have the same restriction [4, 6]).

Coherence optimization: A possible coherence optimization, which we refer to as *delayed-invalidation*, is to delay the processing of invalidation requests. Suppose the local cache C of processor Pi holds a cache line for address a in the shared state. When C receives an invalidation request for a from its parent, C could respond without truly evicting the line, thus letting later loads to read this stale line. However, the stale line must be evicted when C processes a store to a or before Pi commits a Reconcile. A load that reads the stale line is effectively executed early, and the behavior is captured by reading *ib* in WMM. This optimization may violate all dependency orderings, e.g. the behavior in Figure 14. However, if no cache miss can be processed before the eviction of the stale line, then this optimization does not break any dependency ordering or affect the memory model.

Since the stores in WMM are multi-copy atomic (e.g., WMM disallows the behavior in Figure 3), we have

demonstrated that even this coherence optimization is not tied to non-atomic multi-copy stores.

Since WMM allows common coherence optimizations and all speculations in ROB, the performance of its implementation should not be worse than that of any other weak model (e.g. Power and ARM). Furthermore, since Power and ARM cannot directly use load-value prediction, WMM implementation may even have higher performance than Power and ARM.

7. MODELING DATA DEPENDENCY

Figure 14 shows a behavior permitted by WMM but which is not possible unless hardware does load-value prediction or delayed-invalidation optimization. This behavior can be prevented by inserting a Reconcile fence between I_4 and I_5. However, the fence may cause performance loss because it would prevent the execution of loads that follow I_5 but do not depend on I_4. This is an unnecessary cost because commercial microprocessors do not use value prediction yet, and the delayed-invalidation optimization can be adapted to not affect the memory model. To avoid these extra Reconcile fences, we need a memory model that precisely captures the data-dependency ordering enforced in hardware. As we have seen, the axioms of RMO restrict hardware too much, while Power and ARM explicitly models ROB operations. None of these solutions are satisfactory, so we present *WMM-D* which uses *timestamps* to exclude exactly those behaviors that violate data-dependency ordering from WMM.

7.1 Enforcing Data Dependency with Timestamps

We derive our intuition for timestamps by observing how an OOO processor without load-value prediction and delayed-invalidation optimization works. We refer to such a processor as *OOO-D*.

In Figure 14, assume instruction I_k ($k = 1 \ldots 5$) gets its result or writes memory at time t_k in OOO-D. Then $t_5 \geq t_4$ because the result of I_4 is a source operand of I_5 (i.e. the load address). Since I_4 reads the value of I_3 from memory, $t_4 \geq t_3$, and thus $t_5 \geq t_3 \geq t_1$. As we can see, the time ordering reflects enforcement of data dependencies. Thus, a natural way to extend WMM to WMM-D is to attach a timestamp to each value, which will, in turn, impose additional constraints on rule firing in WMM. We first explain how to extend WMM to WMM-D without considering OOO-D, and then show the correspondence between WMM-D and OOO-D.

7.1.1 Adding Timestamps to WMM

Let us assume there is a global clock which is incremented every time a store writes memory. We attach a timestamp to each value in WMM, i.e. an architecture register value, the $\langle$address, value$\rangle$ pair of a store, and a monolithic memory value. The timestamp represents when the value is created. Consider an instruction $r_3 = r_1 + r_2$. The timestamp of the new value in r_3 will be the maximum timestamp of r_1 and r_2. Similarly, the timestamp of the $\langle$address, value$\rangle$ pair of a store (St a v), i.e. the *creation* time of the store, is the maximum timestamp of all source operands to compute $\langle a, v \rangle$. The timestamp of a monolithic memory value is the time when the value becomes visible in memory, i.e. one plus the time when the value is stored.

Next consider a load L (Ld a) on processor i, which reads the value of a store S (St a v). No matter how WMM executes L (e.g. by reading sb, memory, or ib), the timestamp ts of the load value (i.e. the timestamp for the destination register) is always the maximum of (1) the timestamp ats of the address operand, (2) the time rts when processor i executes the last Reconcile fence, and (3) the time vts when S becomes visible to processor i. Both ats and rts are straightforward. As for vts, if S is from another processor j ($j \neq i$), then S is visible after it writes memory, so vts is timestamp of the monolithic memory value written by S. Otherwise, S is visible to processor i after it is created, so vts is the creation time of S.

A constraint for L, which we refer to as *stale-timing*, is that ts should not exceed the time ts_E when S is overwritten in memory. This constraint is only relevant when L reads from ib. In Section 7.1.2, we will explain why this constraint is needed.

To carry out the above timestamp calculus for load L in WMM, we need to associate the monolithic memory $m[a]$ with the creation time of S and the processor that created S, when S updates $m[a]$. When S is overwritten and its $\langle a, v \rangle$ is inserted into $ps[i].ib$, we need to attach the time interval $[vts, ts_E]$ (i.e. the duration that S is visible to processor i) to that $\langle a, v \rangle$ in $ps[i].ib$.

By combining the above timestamp mechanism with the original WMM rules, we have derived WMM-D.

7.1.2 Relation Between WMM-D and OOO-D

The timestamp of each value in WMM-D represents the *earliest* time that the value may become readable in OOO-D. For example, the timestamp of a register value in WMM is the earliest time in OOO-D, at which the value can be derived by an instruction. It should be noted that PC should never be involved in the timestamp mechanism of WMM-D. This is because instructions can be speculatively fetched in OOO-D, and the PC of each instruction can always be known in advance.

As for loads, we first make a simplification that a Reconcile in OOO-D stalls all younger loads, i.e. bypassing from stores younger than the fence is also stalled. This restriction does not reduce the permitted behaviors in OOO-D. This is because no younger load can access memory before the Reconcile commits even without the restriction, and the additionally stalled bypassing can be done immediately after the Reconcile commits.

With the above simplification, a load from processor i can get its value in OOO-D only when (1) its address has resolved, (2) all previous Reconcile fences have committed, and (3) the value is visible to processor i. This exactly corresponds to how we compute the timestamp of the load result in WMM-D. In terms of the time when the value becomes visible in OOO-D, if the load value is also stored by processor i, then the load can bypass

from a store in ROB right after the $\langle$address, value$\rangle$ of the store is computed; otherwise the load must wait for the value to be written into memory. This also corresponds to the computation of vts in WMM-D.

Since it is impossible in OOO-D to have a load get its value after the value has been overwritten in memory, the stale-timing constraint in WMM-D is necessary.

7.2 WMM-D Operational Semantics

The operational semantics of WMM-D is given in Figure 15. We list the things one should remember before reading the rules in the figure.

- The global clock name is gts (initialized as 0), which is incremented when the monolithic memory is updated.
- Each register has a timestamp (initialized as 0) which indicates when the register value was created.
- Each sb entry $\langle a, v\rangle$ has a timestamp, i.e. the creation time of the store that made the entry. Timestamps are added to the method calls on sb as appropriate.
- Each monolithic memory location $m[a]$ is a tuple $\langle v, \langle i, sts\rangle, mts\rangle$ (initialized as $\langle 0, \langle\text{-}, 0\rangle, 0\rangle$), in which v is the memory value, i is the processor that writes the value, sts is the creation time of the store that writes the value, and mts is the timestamp of the memory value (i.e. one plus the time of memory write).
- Each ib entry $\langle a, v\rangle$ has a time interval $[ts_L, ts_U]$, in which ts_L is the time when $\langle a, v\rangle$ becomes visible to the processor of ib, and ts_U is the time when $\langle a, v\rangle$ is overwritten in memory and gets inserted into ib. Thus, the insert method on ib takes the time interval as an additional argument.
- Each processor $ps[i]$ has a timestamp rts (initialized as 0), which records when the latest Reconcile was executed by $ps[i]$.

Some of the timestamp manipulation is done inside the decode and execute methods of each processor $ps[i]$. Therefore we define the following methods:

- decodeTS(): returns a pair $\langle dIns, ts\rangle$, in which $dIns$ is the decoded instruction returned by the original method decode(), and ts is the maximum timestamp of all source registers (excluding PC) of $dIns$.
- executeTS($dIns, ldRes, ts$): first calls the original method execute($dIns, ldRes$), and then writes timestamp ts to the destination register of instruction $dIns$.

We also replace the getRandom method on ib with the following two methods:

- random(a): returns the $\langle$value, time interval$\rangle$ pair of a random stale value for address a in ib. If ib does not contain any stale value for a, $\langle\epsilon, \text{-}\rangle$ is returned.
- rmOlder(a, ts): removes all stale values for address a, which are inserted into ib when $gts < ts$, from ib.

This facilitates the check of the stale-timing constraint.

In Figure 15, WMM-D-Nm and WMM-D-St compute the timestamps of a Nm instruction result and a store $\langle a, v\rangle$ pair from the timestamps of source registers respectively. WMM-D-Rec updates $ps[i].rts$ with the current time because a Reconcile is executed. WMM-D-DeqSb attaches the appropriate time interval to the stale value inserted into ib as described in Section 7.1.1.

WMM-D-Nm rule (Nm execution).

$$\frac{\langle\langle \mathsf{Nm}, dst, v\rangle, ts\rangle = ps[i].\mathsf{decodeTS}()}{ps[i].\mathsf{executeTS}(\langle \mathsf{Nm}, dst, v\rangle, \text{-}, ts)}$$

WMM-D-LdSb rule (Ld execution: bypass from store).

$$\frac{\langle\langle \mathsf{Ld}, a, dst\rangle, ats\rangle = ps[i].\mathsf{decodeTS}() \quad \mathsf{when}(ps[i].sb.\mathsf{exist}(a));\ \langle v, sts\rangle = ps[i].sb.\mathsf{youngest}(a)}{ps[i].\mathsf{executeTS}(\langle \mathsf{Ld}, a, dst\rangle, v, \mathsf{max}(ats, ps[i].rts, sts))}$$

WMM-D-LdMem rule (Ld execution: read memory).

$$\frac{\langle\langle \mathsf{Ld}, a, dst\rangle, ats\rangle = ps[i].\mathsf{decodeTS}();\ \mathsf{when}(\neg ps[i].sb.\mathsf{exist}(a)) \quad \langle v, \langle j, sts\rangle, mts\rangle = m[a];\ vts = (\textbf{if } i \neq j \textbf{ then } mts \textbf{ else } sts)}{ps[i].\mathsf{executeTS}(\langle \mathsf{Ld}, a, dst\rangle, v, \mathsf{max}(ats, ps[i].rts, vts))\ ps[i].ib.\mathsf{rmAddr}(a)}$$

WMM-D-LdIb rule (Ld execution: read stale value).

$$\frac{\langle\langle \mathsf{Ld}, a, dst\rangle, ats\rangle = ps[i].\mathsf{decodeTS}() \quad \langle v, [ts_L, ts_U]\rangle = ps[i].ib.\mathsf{random}(a) \quad \mathsf{when}(\neg ps[i].sb.\mathsf{exist}(a) \ \wedge\ v \neq \epsilon \ \wedge\ ats \leq ts_U)}{ps[i].\mathsf{executeTS}(\langle \mathsf{Ld}, a, dst\rangle, v, \mathsf{max}(ats, ps[i].rts, ts_L))\ ps[i].ib.\mathsf{rmOlder}(a, ts_U)}$$

WMM-D-St rule (St execution).

$$\frac{\langle\langle \mathsf{St}, a, v\rangle, ts\rangle = ps[i].\mathsf{decodeTS}()}{ps[i].\mathsf{executeTS}(\langle \mathsf{St}, a, v\rangle, \text{-}, \text{-})\ ps[i].sb.\mathsf{enq}(a, v, ts);\ ps[i].ib.\mathsf{rmAddr}(a)}$$

WMM-D-Rec rule (Reconcile execution).

$$\frac{\langle\langle \mathsf{Reconcile}\rangle, ts\rangle = ps[i].\mathsf{decodeTS}()}{ps[i].\mathsf{executeTS}(\langle \mathsf{Reconcile}\rangle, \text{-}, \text{-});\ ps[i].ib.\mathsf{clear}();\ ps[i].rts \Leftarrow gts}$$

WMM-D-Com rule (Commit execution).

$$\frac{\langle\langle \mathsf{Commit}\rangle, ts\rangle = ps[i].\mathsf{decodeTS}();\ \mathsf{when}(ps[i].sb.\mathsf{empty}())}{ps[i].\mathsf{executeTS}(\langle \mathsf{Commit}\rangle, \text{-}, \text{-})}$$

WMM-D-DeqSb rule (dequeue store buffer).

$$\frac{a = ps[i].sb.\mathsf{anyAddr}();\ \langle v', \langle i', sts'\rangle, mts\rangle = m[a] \quad ts_U = gts;\ \mathsf{when}(a \neq \epsilon)}{\langle v, sts\rangle \leftarrow ps[i].sb.\mathsf{rmOldest}(a)\ m[a] \Leftarrow \langle v, \langle i, sts\rangle, gts + 1\rangle;\ gts \Leftarrow gts + 1\ \forall j \neq i.\ \textbf{let } ts_L = (\textbf{if } j \neq i' \textbf{ then } mts \textbf{ else } sts')\ \textbf{in}\ \textbf{if } \neg ps[j].sb.\mathsf{exist}(a) \textbf{ then } ps[j].ib.\mathsf{insert}(a, v', [ts_L, ts_U])}$$

Figure 15: WMM-D operational semantics

In all three load execution rules (WMM-D-LdSb, WMM-D-LdMem, and WMM-D-LdIb), the timestamp of the load result is no less than the timestamp of the address operand (ats) or the latest Reconcile execution time ($ps[i].rts$). Besides, the timestamp of the load result is also lower-bounded by the beginning time that the value is readable by the processor of the load ($ps[i]$), In WMM-D-LdSb and WMM-D-LdIb, this beginning time (i.e. sts or ts_L) is stored in the sb or ib entry; while in WMM-D-LdMem, this beginning time is one of the two times (i.e. sts and mts) stored in the monolithic memory location depending on whether the memory value v is written by $ps[i]$ (i.e. whether i is equal to j). In WMM-D-LdIb, the stale-timing constraint requires that $\mathsf{max}(ats, ps[i].rts, ts_L)$ (i.e. the timestamp of the load value) is no greater than ts_U (i.e. the time when the stale value is overwritten). Here we only compare ats with ts_U, because $ts_L \leq ts_U$ is obvious, and the clearing of ib done by Reconcile fences already ensures $ps[i].rts \leq ts_U$.

7.3 Litmus Tests for WMM-D

Enforcing data dependency: First revisit the behavior in Figure 14. In WMM-D, the timestamp of the source operand of I_5 (i.e. the result of I_4) is 2, while the time interval of the stale value 0 for address a in the ib of P1 is $[0, 0]$. Thus I_5 cannot read the stale value 0, and the behavior is forbidden by WMM-D. For a similar reason, WMM-D forbids the behavior in Figure 16, in which I_4 carries data dependency to I_7 transitively. In particular, I_6 reading from I_5 forms a data dependency. This behavior is also impossible in OOO-D.

Allowing other speculations: The behavior in Figure 13 is possible on hardware that performs memory dependency speculation. WMM-D allows this behavior, because the timestamp of the address operand of I_6 is 0, and I_6 can read the stale value 0 from ib. For a similar reason, WMM-D allows the behavior in Figure 4 (assuming I_2 becomes Commit), which can be observed on hardware that speculates over control dependency. As we can see, WMM-D only excludes behaviors that violate data-dependency ordering, while still allowing implementations to speculate on all other dependencies.

Proc. P1	Proc. P2
I_1 : St a 1	$I_4 : r_1 =$ Ld b
I_2 : Commit	$I_5 : r_2 = r_1 + c - 1$
I_3 : St b a	$I_6 : r_3 =$ Ld r_2
	$I_7 : r_4 =$ Ld c
	$I_8 : r_5 = r_4 + a$
	$I_9 : r_6 =$ Ld r_5
WMM-D forbids: $r_1 = a$, $r_2 = a, r_3 = 0$	

Figure 16: Transitive data dependency

Proc. P1	Proc. P2
I_1 : St a 1	$I_4 : r_1 =$ Ld b
I_2 : Commit	$I_5 : r_2 = r_1 + c - 1$
I_3 : St b 1	$I_6 : r_3 =$ Ld r_2
	$I_7 : r_4 =$ Ld c
	$I_8 : r_5 = r_4 + a$
	$I_9 : r_6 =$ Ld r_5
WMM-D allows: $r_1 = 1, r_2 = c$ $r_3 = 0, r_4 = 0, r_5 = a, r_6 = 0$	

Figure 17: RSW in WMM-D

Necessity of two timestamps in memory: In Figure 16, suppose we change I_5 to "St c a", and insert a Commit fence between I_5 and I_6. Then the behavior will be possible in OOO-D, because I_6 can bypass data from I_5, and I_7 can execute early to get 0. (The newly inserted Commit fence cannot stop the bypass). WMM-D also allows the behavior. However, if each monolithic memory location in WMM-D only keeps a single timestamp, which is the time when the memory value becomes visible, then I_6 must get value a from $m[c]$ with timestamp 3. Thus I_7 cannot read stale value 0, which has time interval $[0, 0]$, from ib. This example shows that the two timestamps in each monolithic memory location are indispensable.

Loads to the same address: Remember that two loads to the same address can be executed out of order in OOO-D as long as the two loads read from the same store. WMM-D also captures this subtle optimization. Consider the Read-from-Same-Write (RSW) program in Figure 17. The behavior is observable in OOO-D, because I_7 to I_9 can be executed before I_4 to I_6. It is fine for I_6 and I_7, which read the same address c, to be executed out-of-order, because they both read from the initialization store. WMM-D allows this behavior, because the timestamp of the address operand of I_9 is 0, and I_9 can read stale value 0 from ib. (This behavior is also observable on Power and ARM processors [4, 6]).

In contrast, RMO with the additional RMO-Fix-Ld axiom, which disables the reordering of loads to the same address in Section 4.2, will forbid this behavior (assuming we change I_2 to MEMBAR). This is because the memory order of I_4, I_6, I_7, I_9 must be the same as the program order on P2. This reveals the disadvantage of the axiomatic definition of RMO. Maybe adding complicated axioms can disallow the reordering of loads to the same address while capturing this optimization; we certainly have not figured it out.

7.4 Compiling C++ to WMM-D

The mapping from C++ to WMM-D is almost the same as the one for WMM except that WMM-D does not need any fence for Load Consume in C++. This is because Load Consume leverages data-dependency ordering which is already enforced by WMM-D.

8. MODELING NON-ATOMIC MULTI-COPY STORES

Unlike a multi-copy atomic store, a non-atomic multi-copy store may become visible to different processors at different times. This can happen because of shared store buffers or write-through caches. Even then, all stores for an address can be put in a total order, and the order seen by any processor is consistent with this total order. We will refer to this total order as the *coherence order* ($<_{co}$) [4, 5], though in the literature other names, such as modification order [34], have also been used. We can model such stores by introducing a background rule to make copies of a store in a store buffer into other store buffers. There are quite a few subtleties in doing this properly; *WMM-S* model is an I^{2}E description of rules to generate such behaviors.

8.1 Copying From One Store Buffer into Another

We need a mechanism to identify all the copies of a store in various store buffers. We, therefore, assign a unique tag t when a store is inserted in the store buffer by a store instruction, and this tag is copied when a store is copied from one store buffer into another. When it is time to commit a store from the store buffer to the memory, all the copies of this store have to be deleted from all the store buffers which have them. A store can be committed only if all its copies are the oldest store for that address in their respective store buffers.

All the stores for an address in a store buffer are kept as a strictly ordered list where the *youngest* store (i.e. the *largest* in the total order of this list) is the one that entered the store buffer *last*. We make sure that all ordered lists are can be combined transitively to form a strict partial order, which has now to be understood in terms of the tags on stores because of the copies. By the end of the program, this partial order on the stores for an address becomes the coherence order, so we refer to this partial order as the *partial coherence order*.

Consider the states of store buffers shown in Figure 18. A, B, C and D are different stores to the same ad-

dress, and their tags are t_A, t_B, t_C and t_D, respectively. A' and B' are copies of A and B respectively created by the background copy rule. Ignoring C', the partial coherence order contains:

- $t_D <_{co} t_B <_{co} t_A$ (D is older than B, and B is older than A' in P2) and
- $t_C <_{co} t_B$ (C is older than B' in P3)

Notice that t_D and t_C are not related in this partial order.

At this point, if we allowed C in P3 to be copied as C' into P1, we would introduce a new edge $t_A <_{co} t_C$ in the coherence relation, which would break the partial order by introducing the cycle $t_A <_{co} t_C <_{co} t_B <_{co} t_A$. Therefore copying of C into P1 should not be allowed in this state. Similarly, copying a store with tag t_A into P1 or P2 should be forbidden because it would immediately create the cycle, $t_A <_{co} t_A$. In general, the background copy rule must be constrained so that invariance of the partial coherence order after copying is maintained.

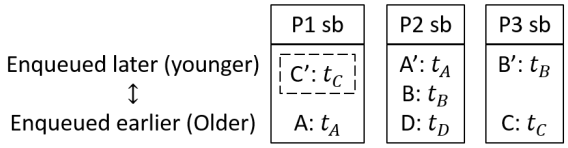

Figure 18: Example states of store buffers by copying stores (primes are copies)

The operational semantics of WMM-S is defined by adding/replacing three rules to the operational semantics of WMM given in Figure 10. These new rules are shown in Figure 19: A new background rule WMM-S-Copy is added to the WMM rules and the WMM-S-St and WMM-S-DeqSb rules replace the WMM-St and WMM-DeqSb rules of WMM, respectively. Before reading these new rules, one should note the following facts:

- The decode method now returns $\langle \mathsf{St}, a, v, t\rangle$ for a store, in which t is the unique tag assigned to the store. Each store buffer entry becomes a tuple $\langle a, v, t\rangle$, in which t is the tag. Tags are also introduced into the methods of sb appropriately.
- The sb now has the following three methods:
  - $\mathsf{has}(t)$: returns True if sb contains a store with tag t.
  - $\mathsf{oldest}(a)$: returns the $\langle$value, tag$\rangle$ pair of the oldest store to address a in sb.
  - $\mathsf{random}(a)$: returns the $\langle$value, tag$\rangle$ pair of any store to address a present in sb.

  If sb does not contain any store to address a, $\mathsf{oldest}(a)$, $\mathsf{random}(a)$ and the previously defined $\mathsf{youngest}(a)$ methods will all return $\langle \epsilon, \epsilon\rangle$.
- A new function $\mathsf{noCycle}(a, t, j)$ is defined to check whether the background rule could copy a store with tag t for address a into the sb of processor j. It returns True if the partial coherence order among the tags of all stores for address a does not contain any cycle after doing the copy.

In Figure 19, WMM-S-St is simply introducing the store tag to the original WMM-St rule. In WMM-S-DeqSb, when we write a store ($\langle a, v, t\rangle$) into memory, we ensure that each copy of this store is the oldest one to address a in its respective store buffers. The old memory value is inserted into the invalidation buffer ib of each processor whose sb does not contain address a. WMM-S-Copy copies a store ($\langle a, v, t\rangle$) from $ps[i]$ to $ps[j]$. The check on $\mathsf{noCycle}(a, t, j)$ guarantees that no cycle is formed in the partial coherence order after the copy. Copying stores from $ps[i]$ to $ps[i]$ will be automatically rejected because noCycle will return False. Since we enqueue a store into $ps[j].sb$, we need to remove all stale values for address a from $ps[j].ib$.

WMM-S-St rule (St execution).

$$\frac{\langle \mathsf{St}, a, v, t\rangle = ps[i].\mathsf{decode}()}{ps[i].\mathsf{execute}(\langle \mathsf{St}, a, v, t\rangle, \text{-})\ \ ps[i].sb.\mathsf{enq}(a, v, t);\ \ ps[i].ib.\mathsf{rmAddr}(a)}$$

WMM-S-DeqSb rule (dequeue store buffer).

$$\frac{\begin{array}{l} a = ps[i].sb.\mathsf{anyAddr}();\ \ old = m[a];\ \ \langle v, t\rangle = ps[i].sb.\mathsf{oldest}(a) \\ \mathsf{when}(a \neq \epsilon \wedge \\ \qquad \forall j.\ \neg ps[j].sb.\mathsf{has}(t)\ \vee\ ps[j].sb.\mathsf{oldest}(a) = \langle v, t\rangle) \end{array}}{\begin{array}{l} m[a] \Leftarrow v \\ \forall j. \quad \textbf{if } ps[j].sb.\mathsf{has}(t) \textbf{ then } ps[j].sb.\mathsf{rmOldest}(a) \\ \qquad \textbf{else if } \neg ps[j].sb.\mathsf{exist}(a) \textbf{ then } ps[j].ib.\mathsf{insert}(a, old) \end{array}}$$

WMM-S-Copy rule (copy store from processor i to j).

$$\frac{\begin{array}{c} \langle v, t\rangle = ps[i].sb.\mathsf{random}(a) \\ \mathsf{when}(t \neq \epsilon\ \wedge\ \mathsf{noCycle}(a, t, j)) \end{array}}{ps[j].sb.\mathsf{enq}(a, v, t);\ \ ps[j].ib.\mathsf{rmAddr}(a)}$$

Figure 19: WMM-S operational semantics

The rule to execute a Commit fence in WMM-S looks the same as that in WMM (i.e. WMM-Com), but has very different implications for implementations. In WMM-S, a store cannot be moved from sb to memory unless all its copies in other store buffers can be moved at the same time. Hence the effect of a Commit fence is no longer local; it implicitly affects all other store buffers/caches. In literature, such fences are known as *cumulative*.

8.2 Litmus Tests for WMM-S

We show by examples that WMM-S allows non-atomic multi-copy store behaviors, and that fences in WMM-S have the cumulative properties similar to those in Power and ARM memory models.

We first revisit the WWC example in Figure 3. The behavior in the figure, which is disallowed by WMM, is now allowed by WMM-S. This is because in WMM-S, I_1 could be copied into the store buffer of P2, and then I_2 reads its value from the store buffer. After that, I_3 is written to memory, I_4 executes, and I_5 is written to memory. Finally I_1 is written to memory, leading the final memory value for a to be 2. This behavior can be found in implementations in which P1 and P2 share a write-through cache. To forbid this behavior in WMM-S, we can insert a Commit fence between I_2 and I_3 on P2 to force I_1 to be written into memory. The inserted Commit fence has a cumulative global effect in ordering

I_1 before I_3 (and hence I_5).

Figure 20 shows another well-known example called Independent Reads of Independent Writes (IRIW). The non-SC behavior in the figure is allowed by WMM-S, while it is forbidden by the original WMM model. This is because in WMM-S, I_1 and I_2 could be copied into the store buffers of P3 and P4 respectively. Then I_3 and I_6 can read the values of I_1 and I_2 from store buffers. After that, I_5 and I_8 simply access the monolithic memory and both get value 0. This behavior can be found in implementations in which P1, P3 share a write-through cache, and P2, P4 share another write-through cache.

To forbid the behavior in Figure 20 in WMM-S, we can insert a Commit fence between I_3 and I_4 on P3, and another Commit fence between I_6 and I_7 on P4. As we can see, a Commit followed by a Reconcile in WMM-S has the same effect as the Power sync fence and the ARM dmb fence. Cumulation is achieved by globally advertising observed stores (Commit) and preventing later loads from reading stale values (Reconcile).

Proc. P1	Proc. P2	Proc. P3	Proc. P4
I_1 : St a 1	I_2 : St b 1	$I_3 : r_1 =$ Ld a	$I_6 : r_3 =$ Ld b
		I_4 : Reconcile	I_7 : Reconcile
		$I_5 : r_2 =$ Ld b	$I_8 : r_4 =$ Ld a
WMM-S allows: $r_1 = 1,\ r_2 = 0,\ r_3 = 1,\ r_4 = 0$			

Figure 20: IRIW in WMM-S

9. CONCLUSION

Weak memory models can be tamed, that is, made more understandable without sacrificing efficiency. One contribution to the complexity is write-through caches, and we see no fundamental advantage of such caches over write-back caches. We do think Instantaneous Instruction Execution (I^2E) descriptions leave little room for ambiguity in the operational semantics of memory models and should be used in all definitions. We have also presented three concrete weak memory models: WMM, a futuristic model that is suitable when load-value prediction becomes commonplace in microarchitectures; WMM-D, a model that enforces data-dependency ordering and can be adopted immediately; and WMM-S, an extension on WMM that allows non-atomic multi-copy store behaviors.

10. ACKNOWLEDGEMENT

This work was done at CSAIL, MIT as part of the Proteus project, which is partially funded by DARPA BRASS grant number 6933274 (2015-2019).

11. REFERENCES

[1] D. Lustig, C. Trippel, M. Pellauer, and M. Martonosi, "Armor: defending against memory consistency model mismatches in heterogeneous architectures," in *Proceedings of the 42nd Annual International Symposium on Computer Architecture*, pp. 388–400, ACM, 2015.

[2] IBM, *Power ISA, Version 2.07*. 2013.

[3] ARM, *ARM Architecture Reference Manual, ARMv7-A and ARMv7-R edition*. 2013.

[4] S. Sarkar, P. Sewell, J. Alglave, L. Maranget, and D. Williams, "Understanding power multiprocessors," in *ACM SIGPLAN Notices*, vol. 46, pp. 175–186, ACM, 2011.

[5] S. Mador-Haim, L. Maranget, S. Sarkar, K. Memarian, J. Alglave, S. Owens, R. Alur, M. M. Martin, P. Sewell, and D. Williams, "An axiomatic memory model for power multiprocessors," in *Computer Aided Verification*, pp. 495–512, Springer, 2012.

[6] S. Flur, K. E. Gray, C. Pulte, S. Sarkar, A. Sezgin, L. Maranget, W. Deacon, and P. Sewell, "Modelling the armv8 architecture, operationally: Concurrency and isa," in *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL 2016, (New York, NY, USA), pp. 608–621, ACM, 2016.

[7] K. Gharachorloo, D. Lenoski, J. Laudon, P. Gibbons, A. Gupta, and J. Hennessy, "Memory consistency and event ordering in scalable shared-memory multiprocessors," in *Proceedings of the 17th International Symposium on Computer Architecture*, pp. 15–26, ACM, 1990.

[8] D. L. Weaver and T. Gremond, *The SPARC architecture manual (Version 9)*. PTR Prentice Hall Englewood Cliffs, NJ 07632, 1994.

[9] L. Lamport, "How to make a multiprocessor computer that correctly executes multiprocess programs," *Computers, IEEE Transactions on*, vol. 100, no. 9, pp. 690–691, 1979.

[10] K. Gharachorloo, A. Gupta, and J. L. Hennessy, "Two techniques to enhance the performance of memory consistency models," in *Proceedings of the 1991 International Conference on Parallel Processing*, pp. 355–364, 1991.

[11] P. Ranganathan, V. S. Pai, and S. V. Adve, "Using speculative retirement and larger instruction windows to narrow the performance gap between memory consistency models," in *Proceedings of the ninth annual ACM symposium on Parallel algorithms and architectures*, pp. 199–210, ACM, 1997.

[12] C. Guiady, B. Falsafi, and T. N. Vijaykumar, "Is sc+ ilp= rc?," in *Computer Architecture, 1999. Proceedings of the 26th International Symposium on*, pp. 162–171, IEEE, 1999.

[13] C. Gniady and B. Falsafi, "Speculative sequential consistency with little custom storage," in *Parallel Architectures and Compilation Techniques, 2002. Proceedings. 2002 International Conference on*, pp. 179–188, IEEE, 2002.

[14] L. Ceze, J. Tuck, P. Montesinos, and J. Torrellas, "Bulksc: bulk enforcement of sequential consistency," in *ACM SIGARCH Computer Architecture News*, vol. 35, pp. 278–289, ACM, 2007.

[15] T. F. Wenisch, A. Ailamaki, B. Falsafi, and A. Moshovos, "Mechanisms for store-wait-free multiprocessors," in *ACM SIGARCH Computer Architecture News*, vol. 35, pp. 266–277, ACM, 2007.

[16] C. Blundell, M. M. Martin, and T. F. Wenisch, "Invisifence: performance-transparent memory ordering in conventional multiprocessors," in *ACM SIGARCH Computer Architecture News*, vol. 37, pp. 233–244, ACM, 2009.

[17] A. Singh, S. Narayanasamy, D. Marino, T. Millstein, and M. Musuvathi, "End-to-end sequential consistency," in *ACM SIGARCH Computer Architecture News*, vol. 40, pp. 524–535, IEEE Computer Society, 2012.

[18] C. Lin, V. Nagarajan, R. Gupta, and B. Rajaram, "Efficient sequential consistency via conflict ordering," in *ACM SIGARCH Computer Architecture News*, vol. 40, pp. 273–286, ACM, 2012.

[19] D. Gope and M. H. Lipasti, "Atomic sc for simple in-order processors," in *High Performance Computer Architecture (HPCA), 2014 IEEE 20th International Symposium on*, pp. 404–415, IEEE, 2014.

[20] SPARC International, Inc., *The SPARC Architecture Manual: Version 8*. Prentice-Hall, Inc., 1992.

[21] J. R. Goodman, *Cache consistency and sequential consistency*. University of Wisconsin-Madison, Computer Sciences Department, 1991.

[22] M. Dubois, C. Scheurich, and F. Briggs, "Memory access buffering in multiprocessors," in *ACM SIGARCH Computer Architecture News*, vol. 14, pp. 434–442, IEEE Computer Society Press, 1986.

[23] X. Shen, Arvind, and L. Rudolph, "Commit-reconcile and fences (crf): A new memory model for architects and compiler writers," in *Computer Architecture, 1999. Proceedings of the 26th International Symposium on*, pp. 150–161, IEEE, 1999.

[24] S. V. Adve and K. Gharachorloo, "Shared memory consistency models: A tutorial," *computer*, vol. 29, no. 12, pp. 66–76, 1996.

[25] L. Maranget, S. Sarkar, and P. Sewell, "A tutorial introduction to the arm and power relaxed memory models." `http://www.cl.cam.ac.uk/~pes20/ppc-supplemental/test7.pdf`, 2012.

[26] J. Alglave, L. Maranget, and M. Tautschnig, "Herding cats: Modelling, simulation, testing, and data mining for weak memory," *ACM Transactions on Programming Languages and Systems (TOPLAS)*, vol. 36, no. 2, p. 7, 2014.

[27] J. Alglave, A. Fox, S. Ishtiaq, M. O. Myreen, S. Sarkar, P. Sewell, and F. Z. Nardelli, "The semantics of power and arm multiprocessor machine code," in *Proceedings of the 4th workshop on Declarative aspects of multicore programming*, pp. 13–24, ACM, 2009.

[28] J. Alglave and L. Maranget, *Computer Aided Verification: 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, ch. Stability in Weak Memory Models, pp. 50–66. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011.

[29] J. Alglave, "A formal hierarchy of weak memory models," *Formal Methods in System Design*, vol. 41, no. 2, pp. 178–210, 2012.

[30] S. Sarkar, K. Memarian, S. Owens, M. Batty, P. Sewell, L. Maranget, J. Alglave, and D. Williams, "Synchronising c/c++ and power," in *ACM SIGPLAN Notices*, vol. 47, pp. 311–322, ACM, 2012.

[31] J. Alglave, D. Kroening, V. Nimal, and M. Tautschnig, "Software verification for weak memory via program transformation," in *Programming Languages and Systems*, pp. 512–532, Springer, 2013.

[32] J. Alglave, L. Maranget, S. Sarkar, and P. Sewell, "Fences in weak memory models (extended version)," *Formal Methods in System Design*, vol. 40, no. 2, pp. 170–205, 2012.

[33] S. V. Adve and M. D. Hill, "Weak ordering a new definition," in *ACM SIGARCH Computer Architecture News*, vol. 18, pp. 2–14, ACM, 1990.

[34] R. Smith, ed., *Working Draft, Standard for Programming Language C++*. `http://open-std.org/JTC1/SC22/WG21/docs/papers/2015/n4527.pdf`, May 2015.

[35] H.-J. Boehm and S. V. Adve, "Foundations of the c++ concurrency memory model," in *ACM SIGPLAN Notices*, vol. 43, pp. 68–78, ACM, 2008.

[36] M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber, "Mathematizing c++ concurrency," in *ACM SIGPLAN Notices*, vol. 46, pp. 55–66, ACM, 2011.

[37] J. Manson, W. Pugh, and S. V. Adve, "The java memory model," in *Proceedings of the 32Nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '05, (New York, NY, USA), pp. 378–391, ACM, 2005.

[38] P. Cenciarelli, A. Knapp, and E. Sibilio, "The java memory model: Operationally, denotationally, axiomatically," in *Programming Languages and Systems*, pp. 331–346, Springer, 2007.

[39] J.-W. Maessen, Arvind, and X. Shen, "Improving the java memory model using crf," *ACM SIGPLAN Notices*, vol. 35, no. 10, pp. 1–12, 2000.

[40] Arvind and J.-W. Maessen, "Memory model = instruction reordering + store atomicity," in *ACM SIGARCH Computer Architecture News*, vol. 34, pp. 29–40, IEEE Computer Society, 2006.

[41] D. J. Sorin, M. D. Hill, and D. A. Wood, "A primer on memory consistency and cache coherence," *Synthesis Lectures on Computer Architecture*, vol. 6, no. 3, pp. 1–212, 2011.

[42] S. Owens, S. Sarkar, and P. Sewell, "A better x86 memory model: x86-tso," in *Theorem Proving in Higher Order Logics*, pp. 391–407, Springer, 2009.

[43] P. Sewell, S. Sarkar, S. Owens, F. Z. Nardelli, and M. O. Myreen, "x86-tso: a rigorous and usable programmer's model for x86 multiprocessors," *Communications of the ACM*, vol. 53, no. 7, pp. 89–97, 2010.

[44] M. Batty, K. Memarian, S. Owens, S. Sarkar, and P. Sewell, "Clarifying and compiling c/c++ concurrency: from c++ 11 to power," in *ACM SIGPLAN Notices*, vol. 47, pp. 509–520, ACM, 2012.

[45] M. Vijayaraghavan, A. Chlipala, Arvind, and N. Dave, *Computer Aided Verification: 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part II*, ch. Modular Deductive Verification of Multiprocessor Hardware Designs, pp. 109–127. Cham: Springer International Publishing, 2015.

[46] G. Z. Chrysos and J. S. Emer, "Memory dependence prediction using store sets," in *ACM SIGARCH Computer Architecture News*, vol. 26, pp. 142–153, IEEE Computer Society, 1998.

[47] M. H. Lipasti, C. B. Wilkerson, and J. P. Shen, "Value locality and load value prediction," *ACM SIGOPS Operating Systems Review*, vol. 30, no. 5, pp. 138–147, 1996.

[48] W. J. Ghandour, H. Akkary, and W. Masri, "The potential of using dynamic information flow analysis in data value prediction," in *Proceedings of the 19th international conference on Parallel architectures and compilation techniques*, pp. 431–442, ACM, 2010.

[49] A. Perais and A. Seznec, "Eole: Paving the way for an effective implementation of value prediction," in *Computer Architecture (ISCA), 2014 ACM/IEEE 41st International Symposium on*, pp. 481–492, IEEE, 2014.

[50] A. Perais and A. Seznec, "Practical data value speculation for future high-end processors," in *International Symposium on High Performance Computer Architecture*, pp. 428–439, 2014.